

MULTIPLE ISP LOCAL AREA NETWORK EGRESS SELECTING METHOD

Field of the Invention

[0001] The present invention relates to the network routing
5 technology, particularly to a method for selecting egresses
of a multi-ISP local area network, and more particularly to
a method for selecting access egresses of a local area network
connected with multiple ISPs.

Background of the Invention

[0002] For convenient description of the present invention and
the prior art, definitions of the following phrases in the
Specification are given hereinafter:

NAT: Network Address Translation;

15 ISP: Internet Service Provider; and

Host route: a corresponding 32-bit mask item of a host
directly connected with a network device in a routing table.
The Address Resolution Protocol (ARP) corresponds to each
32-bit mask item.

20 [0003] In general, a campus network is usually configured with
multiple network egresses to connect with Internet access
providers. The multiple network egresses are backups for each
other and perform load sharing so as to improve the bandwidth
and the robustness for the communication between the campus
25 network and the external network, which is common in school
networks. Usually a school network accesses a public operator

and an education network.

[0004] Furthermore, due to the serious shortage of IP addresses currently, the campus network uses IP addresses in private networks and accesses the Internet through NAT. Since different
5 access providers provide different policies of access accounting and flow control, it is necessary to perform NAT multi-egress policy control on the outgoing traffic of the campus network. Under the NAT multi-egress policy control, it is possible to select an ISP egress in accordance with source
10 information and destination information of subscriber packets to save the charge for access and implement hierarchy management of subscribers. For example, in the school network, the charge for abroad traffic is lower by accessing the public operator (e.g., the Telecom) than by accessing the education
15 network. In contrast, access to the education network for domestic traffic can effectively save charge, since the education network does not charge for domestic traffic. Therefore, it is necessary to perform NAT multi-egress policy control, i.e., an egress of the education network is selected
20 for domestic traffic and an egress of the public operator is selected for abroad traffic, and both of the egresses shall be backups for each other so that all traffic will be switched to one egress when the other fails.

[0005] Currently, there is not a satisfactory solution used in
25 predominant core routing switches for the NAT multi-egress policy control in the industry. Generally, the following solution is adopted in the industry to implement NAT multi-egress policy control:

[0006] Network is built up with multiple routers, one of which performs stream classification on the packets in accordance with source information and destination information of the packets, and the classified packets are transferred to the other routers, which perform NAT operations for the NAT multi-egress policy control.

[0007] FIG.1 is a principle diagram of networking of a multi-ISP campus network in the prior art. It provides the NAT multi-egress policy control function by utilizing a core switch and multiple dedicated NAT devices. This networking can provide the NAT multi-egress policy control function, and the dedicated NAT devices in hardware perform NAT operation to ensure the bandwidth. However, it is required to add the dedicated NAT devices, i.e., each ISP egress needs to be equipped with a dedicated NAT device, resulting in increased cost of networking and increased failure points.

[0008] Furthermore, as described above, the NAT multi-egress policy control function is implemented by the hybrid networking for the multiple devices, so that the core routing switch can not sense the failure of a NAT device when the NAT device at any ISP egress fails, resulting in stream interruption at the ISP egress, and therefore it is necessary to manually modify the complicated stream classification policy to achieve backup of the multiple ISP egresses.

Summary of the Invention

[0009] An object of the present invention aims to provide a method

for selecting egresses of a multi-ISP local area network to solve the problems of increased cost and failure points due to the provision of multiple dedicated NAT devices and to be adapted for the demand of distributed forwarding. It can realize complicated forwarding policy with wire speed and nonblocking, and can also achieve backup for egress links.

[0010] An aspect of the present invention provides a method for selecting egresses of a multi-ISP local area network, including the steps of:

10 providing a distributed NAT board for NAT in the routing switch;

presetting a NAT address pool corresponding to each of the ISP egresses for NAT;

15 querying in a routing table upon request of an outgoing packet from the local area network, and determining the route for the packet; and

determining whether it is necessary to perform NAT at the ISP egress corresponding to the next hop of the route; and if yes, selecting one of the NAT address pools corresponding to the ISP egress, performing corresponding NAT by the NAT board, and forwarding the packet to the egress user board corresponding to the ISP; otherwise, forwarding the packet to the egress user board corresponding to the ISP.

[0011] Preferably, the step of presetting a NAT address pool corresponding to each of the ISP egresses for NAT includes the steps of:

binding each of outgoing interfaces connected with the

ISP with a corresponding one of the NAT address pools; and

creating a NAT policy tree in accordance with combination of the outgoing interface and the source IP address as a keyword upon request for access, wherein leaf nodes of the NAT policy tree store binding relation between each of outgoing
5 interfaces connected with the ISP and the corresponding NAT address pool and the NAT policy information of the slot number of the NAT board.

[0012] Preferably, the step of determining whether it is necessary
10 to perform NAT includes the steps of:

detecting whether there is a public network flag in the routing table item hit by the subscriber traffic; if yes, determining whether one of the leaf nodes of the NAT policy tree is hit in accordance with the combination of the outgoing
15 interface and the source IP address as a keyword; and if one of the leaf nodes of the NAT policy tree is hit, determining it is necessary to perform NAT, otherwise, determining it is unnecessary to perform NAT.

[0013] Preferably, the step of selecting one of the NAT address
20 pools corresponding to the ISP egress includes the steps of:

performing matching in the leaf nodes of the policy tree in accordance with the combination of the outgoing interface and the source IP address as a keyword; and

obtaining the address pool and the slot number of the NAT
25 board from the matched leaf node of the policy tree.

[0014] Preferably, the created NAT policy tree may be a binary tree.

[0015] Preferably, the method may further include the steps of:

classifying the routes of the local area network into a general route and a policy route, and setting a routing policy for the policy route, wherein the general route is a standby for the policy route;

the step of querying in a routing table upon request of an outgoing packet from the local area network and determining a next hop of the route for the packet comprising the steps of:

determining the policy route and/or the general route corresponding to the next hop;

determining whether the ISP egress corresponding to the policy route is available; and if available, replacing the destination address route with the policy routing result; otherwise, utilizing the destination address route of the primary general route.

[0016] Preferably, the step of determining whether the policy route is available includes the steps of:

querying in the routing table in accordance with the next hop of the policy route; and

determining whether the next hop can hit the 32-bit mask route corresponding to a directly-connected host; and if yes, determining the policy route is available, otherwise, determining the policy route is unavailable.

[0017] Preferably, the step of determining a next hop of the route for the packet includes the step of:

determining whether the route corresponds to a plurality

of next hops; and if yes, performing traffic sharing by the plurality of corresponding ISPs.

[0018] Preferably, the routing switch includes a routing module and a NAT module completely separated from each other, wherein
5 the routing module determines route egress for the subscriber traffic; the NAT module determines whether to perform NAT and which NAT address pool to be selected.

[0019] In accordance with the source IP address, the outgoing interface, the general route and the policy router, the
10 embodiments of the present invention determine whether it is necessary to perform NAT for packet forwarding and determine the binding relation between the address pools and the outgoing interfaces, and add the NAT strategy tree describing the binding relation with the address pool. Moreover, the routing
15 module is separated from the NAT module so as to meet the demand of distributed forwarding, implementing complicated forwarding policy and wire speed and nonblocking forwarding.

Brief Description of the Drawings

20 [0020] FIG.1 is a networking principle diagram of a conventional multi-ISP campus network;

[0021] FIG.2 is a flow diagram of a NAT multi-ISP policy forwarding according to an embodiment of the present invention; and

[0022] FIG.3 is a flow diagram showing the NAT policy forwarding
25 with reference to a particular device according to the other embodiment of the present invention.

Detailed Description of the Embodiments

[0023] The key of the preferred embodiments of the present invention lies in that a NAT policy table is added in the forwarding plane where NAT policy control is directly performed when forwarding the data stream, so that the core routing switch can not only accomplish the complicated NAT policy control, but also take advantage of high performance of the distributed forwarding plane. Thus, the core routing switch at the core position of the campus network can independently provide the NAT policy function, resulting in simplification of the network.

[0024] The following two major technical problems can be solved by addition of the NAT policy table to realize the NAT policy function:

[0025] 1. It enables selection of ISP egresses in accordance with subscriber source information and outgoing interface information, flexible accounting and flow control policy can be implemented in combination with multiple ISPs, and the charge for outgoing traffic of subscribers can be saved.

[0026] 2. It enables hot backup among the multiple ISP egresses, i.e., automatically and rapidly switching to another ISP egress without any manual intervention upon detection of a failed ISP egress.

25 [0027] FIG.2 is a flow diagram of NAT multi-ISP policy forwarding according to a preferred embodiment of the present invention. The NAT policy table is added in the forwarding plane. The

policy table is stored in a tree form. The index of the table is the source IP address plus the outgoing interface, and the content in the items of the table is bound ISP egress information including address pool, restriction on the number of links, etc. During the forwarding, query is performed in the routing table and the policy route, and performs query in the NAT policy table, obtains ISP egress information and performs the NAT according to the obtained ISP egress information. When the ISP egress fails, selecting an available ISP automatically to achieve hot backup among the multiple ISPs.

[0028] The forwarding steps will be described in detail as follows:

[0029] 1. A forwarding outgoing interface A is determined for a packet by querying in the routing table in accordance with the destination IP address of the packet;

[0030] 2. The flow determines whether it is necessary to perform policy routing in accordance with system configuration information, and if unnecessary, the flow performs query in the NAT policy tree by using the source IP address plus the outgoing interface A, and then jumps to step 5; if necessary, the flow executes step 3;

[0031] 3. An outgoing interface B is determined for the packet by performing policy routing in accordance with the result of complicated stream classification;

[0032] 4. The flow determines whether the outgoing interface B is valid, and if the outgoing interface B is valid, the flow

performs query in the policy tree by using the source IP address plus the outgoing interface B; if the outgoing interface B is invalid, the flow performs query in the policy tree by using the source IP address plus the outgoing interface A; and

5 [0033] 5. An ISP egress is selected in accordance with the query result of the NAT policy tree, the NAT operation is performed on the packet, and the packet is sent out over the link corresponding to selected ISP egress.

[0034] The backup of the multiple ISP egresses may be implemented
10 by the following two means:

[0035] 1. As for the policy route, an outgoing interface of the general route will be utilized automatically if the outgoing interface of the policy route is invalid.

[0036] 2. As for the general route, if the outgoing interface of
15 the general route is invalid, the route processing system of the core routing switch will automatically perform route recalculation, select a new route, and distribute the new route in the routing table, so as to achieve the backup of multiple ISP egresses.

20 [0037] In order to meet the controllability requirement on devices in the campus network connected with multiple ISP egresses in case of hybrid networking with multiple address spaces, the policy NAT in the preferred embodiment of the present invention will realize the following three critical functions:

25 [0038] A. The outgoing egress for subscriber traffic should not be determined merely by the general route. The improved policy route has to be completed and the backup must be implemented

- 11 -

for the policy route through the general route.

[0039] B. The following requirements must be met: subscribers within one private network can access the public network via egresses provided by different ISPs; and when one subscriber
5 in the private network accesses the public network via different egresses, the address of the subscriber in the private network can be translated into an address in the public network in different address pools, i.e., the NAT must be performed by the address pool bound with the egress when the
10 address space of the subscriber is not consistent with his egress space.

[0040] C. The routing module is completely separated from the NAT module: the routing module (including the destination address route and the policy route) determines the egress of the
15 subscriber traffic, and the NAT module determines whether to perform NAT and which address pool to select.

[0041] In order to attain the object that the egress of the subscriber traffic should not be determined merely in accordance with the general route, the improved policy route
20 must be completed, and the backup must be implemented for the policy route through the general route, the general route is utilized to backup the policy route, i.e., the subscriber traffic will be forwarded automatically in accordance with the general route when the policy route is not available. In the
25 embodiment of the present invention, the next hop of the policy route is searched in the routing table, since the next hop of the available route generally corresponds to a directly-connected host, whether the 32-bit mask route

corresponding to the directly-connected host can be hit will be taken as the criterion for determining whether the policy route is available. If the 32-bit mask route corresponding to the directly-connected host can be hit, the policy route is available; if the 32-bit mask route corresponding to the directly-connected host can not be hit, the policy route is invalid and thus the general route is utilized for forwarding.

[0042] In order to implement that subscribers in one private network can access the public network via the egresses provided by different ISPs, and that when a subscriber in one private network accesses the public network via different egresses, the address of the subscriber in the private network shall be translated into an address in the public network in different address pools, i.e., the NAT must be performed by the address pool bound with the egress for the subscriber traffic when the address space of the subscriber is not consistent with his egress space, thus the embodiment of the present invention performs translation on different subscriber traffic in accordance with different address pools via different ISP egresses in the embodiment of the present invention. In the embodiment of the present invention, the address pools are not registered in global mode but bound with outgoing interfaces. Meanwhile, in order to identify whether it is necessary to perform NAT and which NAT address pool to be selected, a NAT policy tree is created through combination of the outgoing interface and the source IP address, recording the binding relation of the address pools and the slot number of the distributed dedicated NAT board. Whether there is a public network flag in the routing table item hit by the subscriber

traffic is taken as the enablement switch to search in the NAT policy tree. The flag is configured by the subscriber at the outgoing interface connected with an ISP, and any route related with the outgoing interface contains such a public network flag.

5 If using the combination of the outgoing interface and the source IP address as a keyword can hit a leaf of the NAT policy tree, it indicates that it is necessary to perform NAT before the packet is sent out, thus the address pool and the slot number of the NAT board are obtained from the leaf of the NAT policy
10 tree, and the packet is forwarded to the NAT board to process; otherwise, it indicates that the address of the subscriber is an address in the public network, thus the subscriber and the ISP connected herewith pertain to the same address space, therefore, it is unnecessary to perform NAT, and the packet
15 is forwarded to a corresponding ISP egress subscriber board to process in accordance with the route information.

[0043] For complete separation of the routing module from the NAT module, the routing module (including the destination address route and the policy route) determines an egress for the
20 subscriber traffic, and the NAT module determines whether to perform NAT and which address pool to be selected. The embodiment of the present invention adopts the complete separation of the routing module from the NAT module to ensure clear logical separation and no influence in function between
25 them, so that there is sufficient space for achieving combination of forwarding logics of various complicated streams from different subscribers.

[0044] FIG.3 is a flow diagram showing the NAT policy forwarding

- 14 -

with reference to a particular device according to another preferred embodiment of the present invention.

[0045] In step 210, the flow performs searching in the routing table in accordance with the destination IP address, to
5 determine a possible next hop in accordance with the routing table;

[0046] In step 220, the flow determines whether there are multiple next hops according to the searched routing table;

[0047] If there are multiple next hops in step 220, the flow
10 performs traffic sharing on the multiple next hops in step 230, and then goes to step 240, where the flow determines whether the policy route is matched successfully;

[0048] If the flow determines in step 220 that there is no multiple next hops but only one next hop, the flow will directly go to
15 step 240, where the flow determines whether the policy route is matched successfully;

[0049] If it is determined in step 240 that there is a successful match for the policy route, the flow goes to step 250, where it is determined, by searching in the routing table in
20 accordance with the next hop of the policy route, whether the route of a host can be hit; if it is determined in step 250 that the route of a host can be hit, the flow will go to step 260, where the destination address of the route is covered with the searching result in the policy route. Then the flow goes
25 to step 270, where whether there is a public network flag in the routing table item is determined;

[0050] If it is determined in step 240 that there is an

unsuccessful match for the policy route, or if it is determined in step 250 that the route of a host can not be hit, the flow will go to step 270;

[0051] If it is determined in step 270 that there is a public
5 network flag, goes to step 280, where whether a leaf of the NAT policy tree is hit is determined by searching in the NAT policy tree in accordance with the source IP address and the outgoing interface;

[0052] If it is determined in step 280 that a leaf of the NAT policy
10 tree is hit, the flow goes to step 290, where an address pool number is obtained in accordance with the searching result. The packet is in turn forwarded in step 310 via the switching network to the distributed NAT processing device of the NAT board to perform NAT. In step 300, the packet will enter the
15 switching network;

[0053] If it is determined in step 270 that there is no public network flag, or if it is determined in step 280 that the route of a host is not hit, the flow will go to step 300 to process via the switching network;

20 [0054] Finally, in step 320, the packet is forward to the egress user board in accordance with the routing result.

[0055] The above descriptions are preferred embodiments of the present invention, wherein the described methods are merely for the purpose of exemplification, and not intended to limit
25 the scope claimed for the invention, and all the equivalent variations of the description and the appended drawings shall be included in the scope of claims of the present invention.